Energy-Efficient System Architecture Aims to Improve System Energy Efficiencies

Ravi Nagaraj

Table of Contents

(Click on page number to jump to sections)

Energy-Efficient System Architecture Aims to Improve System Energy Efficiencies

Overview: Reducing Power Consumption

Minimizing power consumption is an objective most computing users have. For laptop users, it affects battery life. For desktop users, it affects cooling needs and therefore can impact distracting fan noise in an office. For server users, the cost of power and cooling are a growing part of data center operating expenses. To help address these challenges, researchers in the Systems Technology Lab, part of the Intel Corporate Technology Group, are focused on establishing a new computing architecture based on power management. Intel's goal is to be the leader in computing performance per watt.

Instead of looking at small architectural changes to achieve incremental power savings, Intel researchers envision an all-new Energy-Efficient System Architecture (EESA) designed to increase performance per watt by managing voltage and frequency. EESA uses sensors to identify when a system function is idle and then sends that component to its lowest state of power consumption.

EESA is expected to result in world-class energy efficiency that will extend Intel innovation and leadership in computing performance per watt. Intel is inviting developers from the PC industry to help define specifications and interfaces in the areas of I/O device optimization, system power conversion, and sensor architecture.

Five Key Technology Areas Comprise EESA

The developers of EESA envision five technologies working together to optimize energy efficiency. Those functions are:

- Fine-Grain Power Management
- I/O Optimization
- System Power Conversion
- Client Sensor Architecture
- Power Policy Management

Fine-Grain Power Management (FGPM)

The *Fine-Grain Power Management* (FGPM) function is central to EESA and provides more precise control over power levels inside the system. Today, the system sleep function is initiated only when all components are idle. As shown in **Figure 1,** FGPM, an Intel technology, will put individual functions of the system to sleep when they are idle, resulting in greater power savings.

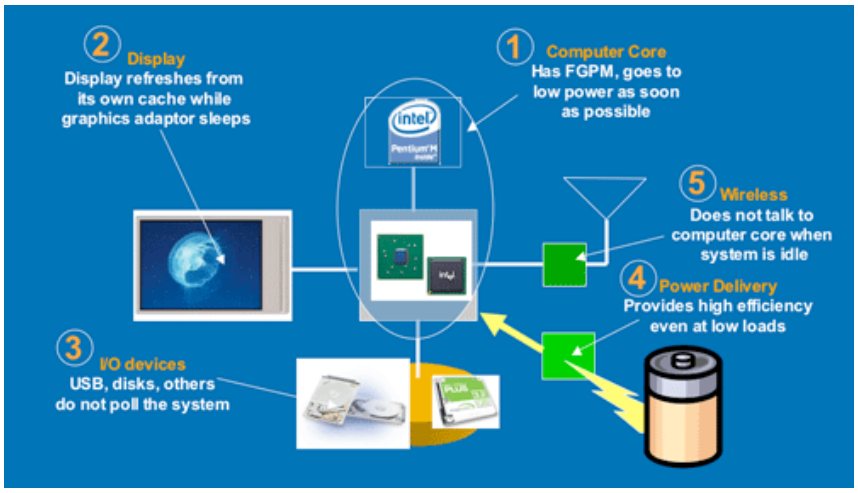

Figure 1. Using sensors built into the chipset, FGPM will send idle parts of the system into a low power state as soon as possible.

I/O Optimization

I/O Optimization technology can lower power consumption in three user-facing areas: self-refreshing displays, power-managed I/O, and self-refreshing audio. I/O Optimization technology looks at the interfaces to these devices to find ways to reduce their dependency on processing cycles from the system. If system functions can be put into idle state more often, power consumption will be lowered.

Self-refresh displays take advantage of the fact that the content of most computer displays does not change very rapidly, whether it is associated with a desktop, laptop, PDA, or cell phone. It is the nature of human interactions with computers that it takes users awhile to assimilate the information being displayed; therefore displays can remain unchanged for a few seconds to many minutes at a time. Meanwhile, the graphics logic behind the display is consuming power by refreshing the screen many times per second. Intel researchers invented the self-refresh displays function to enable that display logic to be put to sleep when the screen is not being updated. The idea is to install a small amount of memory—which is increasingly inexpensive—in the display panel itself. This cache memory stores the information to be displayed to the user. During the time when the display is not changing, the display gets its information from that cache and power is reduced by I/O Optimization technology putting the rest of the display logic to sleep.

In the future, self-refresh display technology will be incorporated into mobile devices. When a wireless network polls the device, an EESA device will respond from a self-refresh cache rather than waking the entire processor and chipset to respond to the poll.

Power-managed I/O technology provides feedback to the FGPM function about the power management state of Universal Serial Bus (USB) devices built into the computer (external I/O devices are not affected). If the USB device is idle, FGPM puts it, and the chipset driving it, to sleep until their services are needed.

The *self-refreshing audio* function maintains a large buffer of audio that is used to drive audio output to the user. Then, the system disk drive is accessed only occasionally when the buffer needs to be refreshed. In between those accesses, the disk, the processor and other unused circuitry can be set at low power or in energy conservation state. The user continues to hear uninterrupted audio, yet the rest of the machine is powered down for enough time to have an affect on energy consumption.

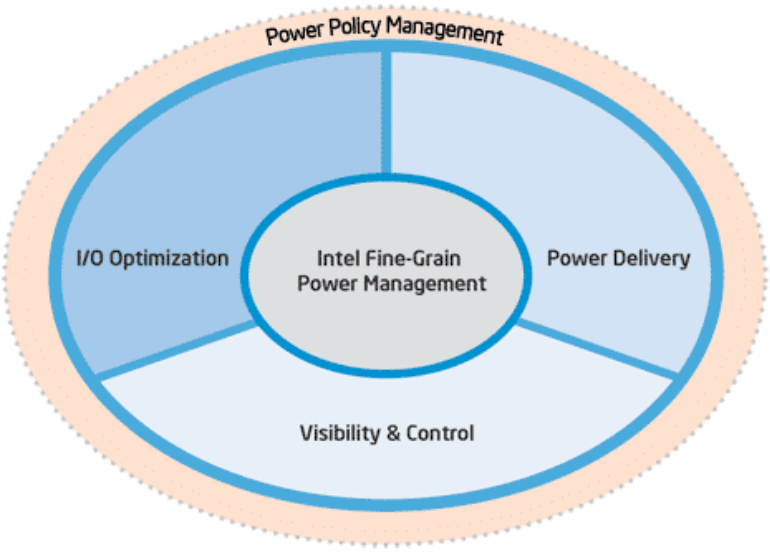

Figure 2. *I/O Optimization and EESA's three other function areas operate within the Power Policy Management technology's operating parameters.*

System Power Conversion

The *System Power Conversion* function streamlines how power is managed from its source to the circuits that use it. Repeated power conversions in today's computers result in delivery efficiencies as low as 50 percent—meaning that half the power is consumed in conversion processes themselves. Intel researchers are developing methods to improve that efficiency, with 90 percent as their goal.

In addition, the System Power Conversion function would react to increased system temperatures by slowing the clock rate for some components to increase power conversion efficiency. Slowing the clock reduces the power consumed in the CMOS circuits, lowers the heat and allows increasing power converter efficiency. For blade servers, sensors could feed information to FGPM to help balance compute loads across a rack of blades, thereby reducing thermal output for the rack as a whole. For example, if one blade is working hard and generating a lot of heat, a second underutilized blade could take some of that work so the total heat generated by both blades is reduced. In addition, the System Power Conversion function will allow idle components of a blade server to switch off to conserve power and reduce heat generation. That reduced need for cooling could further reduce power consumption when some fans are also shut down.

These same System Power Conversion principles can also benefit handheld devices. System Power Conversion technology can convert battery power more efficiently—giving mobile users more computing time.

Client Sensor Architecture

Client Sensor Architecture standardizes how sensors communicate back to FGPM. Today there is no mechanism for many sensors to report their information back to a central management function. Client Sensor Architecture is designed to organize sensors by identifying their capabilities, monitoring their functions, and defining an orderly process to get information to FGPM in a systematic way.

Power Policy Management

Power Policy Management technology, incorporates all the above information and controls, as shown in **Figure 2**, to maximize overall power efficiency for the platform. More information about the component parts of EESA will be released over the next several IDF cycles.

Summary

Energy-Efficient System Architecture combines five new technologies to improve computing performance per watt for all types of computing devices. Improvements will be made by managing voltage and frequency to reduce power demands within the chipset.

At the Spring 2006 Intel Developer Forum (IDF), attendees were given an EESA overview and asked to join in the development effort. Developers should expect updates at the upcoming Fall 2006 and Spring 2007 IDFs.

For portable devices, EESA will boost power conversion efficiencies by up to 80 percent, resulting in a longer battery life. For developers, EESA presents challenges in building each of the five functions and integrating them into a single system architecture to overcome future power and thermal barriers in chipset design.

More Info

Learn more about how Intel is bringing the benefits of energy-efficient performance to the world at the Intel Web site.

Author Bio

Ravi Nagaraj
Principal Engineer
System Technology Lab
Corporate Technology Group
Intel Corporation

Ravi Nagaraj manages the hardware aspects of the Low Power IA group in the Corporate Technology Group at Intel. He joined Intel in 1989, and has held various positions there in system design and architecture of mobile PCs. From 1999 through 2000, Nagaraj briefly left Intel to become director of engineering at ViA, Inc., a start-up developing wearable computers in Minnesota. Nagaraj's current focus is architecture research to reduce the system power of Intel® architecture-based small-form-factor devices. He graduated summa cum laude from Penn State University with an M.S.E.E. He has been awarded 15 U.S. patents and has filed several others.

—End of Technology@Intel Magazine Article—